FOR FURTHER TRAN

12

ARI TECHNICAL REPORT
TR-77-A26

AD A055657

Application of Adaptive Decision Aiding Systems to Computer-Assisted Instruction

by

Donald M. May, William H. Crooks, Denis D. Purcell,
Luigi F. Lucaccini, Amos Freedy, and Gershon Weltman

PERCEPTRONICS
6271 Variel Avenue
Woodland Hills, California 91367

OCTOBER 1977

Contract DAHC 19-74-C-0027

Prepared for

U.S. ARMY RESEARCH INSTITUTE
for the BEHAVIORAL and SOCIAL SCIENCES
5001 Eisenhower Avenue
Alexandria, Virginia 22333

U. S. ARMY RESEARCH INSTITUTE

FOR THE BEHAVIORAL AND SOCIAL SCIENCES

A Field Operating Agency under the Jurisdiction of the Deputy Chief of Staff for Personnel

J. E. UHLANER
Technical Director

W. C. MAUS
COL, GS
Commander

Research accomplished
under contract to the Department of the Army

PERCEPTRONICS, INC.

NOTICES

ARI TR-77-A26

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER
Technical Report TR-77-A26

2. GOVT ACCESSION NO.

3. RECIPIENT'S CATALOG NUMBER

4. TITLE (and Subtitle)
6
APPLICATION OF ADAPTIVE DECISION AIDING SYSTEMS TO COMPUTER-ASSISTED INSTRUCTION.

5. TYPE OF REPORT & PERIOD COVERED
9
Final Report.
Jan[redacted] - Dec[redacted]74

6. PERFORMING ORG. REPORT NUMBER
14
PATR-1019-74-12

7. AUTHOR(s)
10
Donald M./May, William H./Crooks, Denis D./Purcell, Luigi F./Lucaccini, Amos/Freedy, Gershon Weltman

8. CONTRACT OR GRANT NUMBER(s)
15
DAHC19-74-C-0027 new

9. PERFORMING ORGANIZATION NAME AND ADDRESS
Perceptronics, Inc.
6271 Variel Avenue
Woodland Hills, California 91364

10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
16
2Q161102B74B

11. CONTROLLING OFFICE NAME AND ADDRESS
US Army Research Institute for the Behavioral and Social Sciences PERI-OK
5001 Eisenhower Avenue, Alexandria, VA 22333

12. REPORT DATE
11
Oct[redacted]1977

13. NUMBER OF PAGES
46
12 55p.

14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)
18 ARI / 19 TR-77-A26

15. SECURITY CLASS. (of this report)
UNCLASSIFIED

15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for Public Release; Distribution Unlimited

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

Research was monitored technically by Leon H. Nawrocki and Bruce W. Knerr of the Educational Technology and Training Simulation Technical Area, ARI.

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

COMPUTER-ASSISTED INSTRUCTION	ELECTRONIC TROUBLESHOOTING	DIAGNOSIS
ADAPTIVE LEARNING SYSTEMS	DECISION TRAINING	TRAINING
COMPUTER-AIDED DECISIONS	ADAPTIVE CAI	
ARTIFICIAL INTELLIGENCE	TRAINABLE MACHINE	
TRAINING TECHNOLOGY	ADAPTIVE DECISION AIDING	

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

A minicomputer-based Computerized Diagnostic and Decision Training (CDDT) system, combining principles of artificial intelligence, decision theory, and adaptive computer-assisted instruction, is described. Training focuses on decision-making in electronic troubleshooting. The CDDT System incorporates an adaptive computer program which learns the student's diagnostic and decision value structure using a trainable network technique of pattern classification, compares this structure to that of an expert, and adapts the instructional sequence to eliminate discrepancies through the use of heuristic algorithms. (over)

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

78 06 22 015

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

390 285

JOB

20.

An expected value model of decision-making is the basis of the student and instructor models which, with the task simulator and adaptive instructions, form the core of the CDDT system. The instructor model also generates suggested actions in response to student requests for assistance. The student's task is to troubleshoot a complex circuit by making test measurements, replacing the malfunctioning part, and making verification measurements. The student values of interest are those for information gained through the measurements, and for the replacement of circuit modules.

FOREWORD

The Educational Concepts and Evaluation Work Unit Area of the Army Research Institute for the Behavioral and Social Sciences (ARI) performs research in areas of educational technology which may be applicable to military training. Of special interest are development and implementation of computer-based training systems to ease such current Army problems as a shortage of qualified instructor personnel, a student population of widely varying abilities, and increased training costs.

Training is more effective if instruction is adapted to the individual student. Programmed Instruction (PI), tutorial Computer-Assisted Instruction (CAI), and Computer-Managed Instruction (CMI) are methods by which instruction can be individualized. PI only allows a student to proceed through a fixed instructional sequence at his own pace. CAI and CMI can theoretically provide unlimited amounts of individualization, but they are, in practice, severely limited by the programming resources required to provide such individualization. The present Technical Report describes the first phase of a research effort to develop a technique for overcoming this limitation through the use of artificial intelligence (AI) techniques.

Artificial intelligence theory and techniques constitute an approach having great potential benefit for CAI system development. Among the advantages of AI--computer programs which exhibit "intelligent" behavior--is the opportunity for mixed-initiative interaction between student and computer.

The effort was part of the ARI technological base program, which investigates areas where progress in resolving critical Army problems has been inhibited by a lack of understanding of fundamentals or a scarcity of basic data. ARI provided guidance and technical monitoring to the work done under Contract DAHC 19-74-C-0027 by Perceptronics, Inc., an organization selected as having unique capabilities for research and development in this area.

The research is responsive to the requirements of Army Project 2Q161102B74B and to special requirements of the Product Manager, Computerized Training Systems.

APPLICATION OF ADAPTIVE DECISION AIDING SYSTEMS TO COMPUTER-ASSISTED INSTRUCTION

BRIEF

Requirement:

To design and develop an adaptive computerized training system using "artificial intelligence" techniques; to develop the necessary instructional materials and sequence to apply the system to an electronic troubleshooting task; and to install the system on a minicomputer and demonstrate its operation.

Research Product:

A previously developed program with an adaptive decision-making capability was modified to operate in a training context. The necessary simulated electronic circuit, adaptive instructions, input-output routines, and supporting software were developed and installed.

The system incorporates an adaptive computer program which learns the student's diagnostic and decision value structure, compares this structure to that of an expert, and changes the instructional sequence (by providing feedback and new problems) to modify the student's value structure until it matches that of the expert. An expected value model of decision-making is the basis of the student and instructor (expert) models, which, in conjunction with the task simulator and adaptive instructions, form the core of the training system. The student model is dynamically adjusted using a trainable network technique of pattern classification. Heuristic algorithms generate the adaptive instructions and modify the problem presentation sequence. The instructor model also generates suggested actions in response to student requests for assistance. The student's specific task was to troubleshoot a simulated complex electronic circuit by making various test measurements, replacing the malfunctioning part, and making final verification measurements. The student values of interest are those for information gained through the measurements, and for replacement of circuit modules.

Utilization of Findings:

The working system demonstrates the feasibility of applying artificial intelligence techniques to computer-assisted instruction in a minicomputer environment. On a purely theoretical basis, the system appears to have potential application to Army training. However, any operational implementation must be dependent on future research investigating the training effectiveness of such a system.

CONTENTS

CONTENTS (Continued)

FIGURES

1. ADAPTIVE TECHNIQUES IN COMPUTER-ASSISTED MAINTENANCE TRAINING

The use of computers for instruction has increased in recent years, yet despite increasing sophistication of both computer hardware and software, there have not been corresponding advances in techniques for the use of computers in the instructional process. Many of the instructional concepts and techniques currently implemented in Computer-Assisted Instruction (CAI) originated prior to the advent of CAI. However, as new techniques in computer science, and in particular in the field of Artificial Intelligence (AI), become available, it will be possible to implement more of them within CAI systems. Such techniques appear to have the potential to provide individualized instruction to a greater extent than is currently possible. This section provides a review of previous research in the areas of CAI, AI, and adaptive decision models, which forms the basis for the development of the Computerized Diagnostic and Decision Training (CDDT) system.

CAI and Individualized Instruction

Goals of CAI

A central theme in the field of educational technology is the creation of methods which allow the individualization of instruction. Training specialists and educational theorists recognize the importance of focusing on the individual student if significant advances in the efficiency and effectiveness of instruction are to be made (e.g., Crawford & Ragsdale, 1969; Glaser, 1965). Bloom (1968) has advocated the concept of mastery learning, in which instruction is designed and managed so that all students reach a given level of achievement, albeit at different rates.

The principles now included under the rubric of Programmed Instruction (PI), which grew out of pioneering work by Pressey, Skinner, and others, have facilitated the practical implementation of mastery learning techniques. Such principles, also claimed as advantages of PI, include student-paced progression, immediate knowledge-of-results, individualized instructional sequencing, use of explicit performance objectives, diagnostic assessment, and the division of instruction into small discrete steps. These principles formed the basis for the multiplicity of programmed textbooks, teaching machines, and early CAI systems seen in the 1960's.

Development of CAI Systems

Much of the early work in CAI included the direct incorporation of PI techniques on computers, with the objective of developing fast, individualized instructional systems which would relieve the human teacher from the burden of routine drill, tutoring, and instructional bookkeeping. This early enthusiasm for PI techniques applied in the

computer environment is documented by the exhaustive work published by the National Education Association (Lumsdaine & Glaser, 1960). Grubb (1971) illustrated the full range of possible instructional bookkeeping across a variety of CAI modes.

Experiments at Stanford University were among the first large-scale attempts to develop models for optimizing the instructional process in CAI drill and practice situations (Atkinson & Wilson, 1968; Groen & Atkinson, 1966; Suppes & Morningstar, 1969). An alternative to the prescriptive approach of the Stanford studies is provided by the learner-controlled method (Grubb, 1969). The Stanford approach employs mathematical models to optimize presentation strategies in accordance with the student's response history; the learner-controlled method is based on the assumption that prescription of instruction is not yet possible and CAI systems must therefore possess the flexibility for students to select their own unique instructional sequences. Although these approaches are radically different, they both illustrate the growing recognition that instructional systems must adapt to individual student needs, abilities, and interests.

CAI Techniques

Researchers have divided CAI into as many as a dozen categories (Zinn, 1970). For this review it is useful to define four major categories; these are (a) Simulation; (b) Tutorial; (c) Drill; and (d) Inquiry. These modes of CAI will be discussed below.

Simulation. "Simulation" and "gaming" are often used interchangeably, although a "game" often connotes a degree of direct competition. One way of viewing a simulation is that it is a computerized version of a simulator--that is, a manageable representation of a real system or task. Simulation has been used for years in industry and in the military as a practical aid to decision makers. The simulation models a dynamic real world environment in which the operator makes decisions. The outcome of these decisions is then simulated and the result is displayed.

In CAI, simulation is usually combined with other CAI modes as well as audio-visual aids and textbooks. Laboratory assignments are simulated in chemistry, physics, medicine, and electronics. This eliminates the need for costly laboratory facilities and permits experience in simulated research laboratories which could not possibly be made available for instructional purposes. This has the effect of permitting small instructional facilities to offer courses of the same high quality as those offered by larger facilities. In problems requiring a diagnosis, the instructor can easily input the parameters of a malfunction, while the evaluation software accompanying a simulation package can provide evaluations more detailed than the teacher can produce.

Tutorial. The tutorial CAI program presents instructional material to the student and then asks the student questions which are usually of

the selected response type. The student indicates an answer; and sometimes the CAI system also requests the student's level of confidence in the correctness of the answer (Shuford, 1965). Based on the student's response, and often on other criteria (such as response history, pre-instruction test results, level of confidence or indicated preference), a new instructional sequence is selected for the student.

A good tutorial system will have a rich branching structure and many decision rules. With tutorial CAI the program designers (i.e., teachers and programmers) must have a clear concept of their teaching philosophy and instructional strategies, and must know a great deal about the system. Tutorial programs often employ drill, simulation, inquiry, and dialogue modes as well.

Drill. In the drill mode of CAI, the student is presented with problems to be worked or questions to be answered. The system presents additional problems or questions based on the student's response. This mode is easy to design. Instructional material is presented in the classroom or by another CAI mode. At Stanford individualization is achieved through off-line update based on overnight evaluation of student performande and selection of appropriate lesson material (Atkinson & Wilson, 1968).

Inquiry. In the inquiry mode of CAI, the student can ask questions for which the program has stored answers. This requires efficient information storage structures and searching algorithms. The inquiry mode can be extended to what is known as the dialogue mode in which the computer may respond to the student's question (or response) with a suggestion, another question, or an instructional prescription. Dialogue programs are very difficult to construct because a vocabulary and decision response structure must be designed which will meet most contingencies.

Adaptive CAI

Previous Approaches

It has been recognized for more than a decade that true individualized instruction must include some form of adaptation to the individual student (Smallwood, 1962). However, while most researchers recognize the need to adapt instruction to individual differences, adaptation is usually made on the basis of response history. That is, the great majority of adaptive programs are made adaptive by the branching structure of the programs rather than by the use of Artificial Intelligence techniques. Various criteria are used as a basis for the adaptive process. Some of these are: the student's performance on the learning task (Melaragno, 1966; Smallwood, 1962); the student's level of confidence in correctness of the response (Shuford, 1965; Shuford & Massengill, 1967; Baker, 1965; Kopstein & Seidel, 1969); learning time; the past stimulus-response history of the student (Groen & Atkinson, 1966; Smallwood, 1962); and the fitted parameters of a learning model (Atkinson, 1972).

Similarly, a variety of system outputs have been made to adapt; these include: the content of the presented information (Stolurow, 1969); the presentation sequence (Stolurow, 1969; Kopstein & Seidel, 1969); the presentation rate (Stolurow, 1969; Kopstein & Seidel, 1969); the nature and pacing of reinforcement (Groen & Atkinson, 1966); the choice of the next presentation (Groen & Atkinson, 1966); and the amount of time spent on the computer terminal (Atkinson, 1972).

In addition to the applications of adaptive training in CAI systems mentioned above, a number of adaptive techniques from control theory have been introduced to training. This is particularly evident in perceptual-motor skill training (Kelley, 1969; Lowes, Ellis, Norman & Matheny, 1968). In these implementations of adaptive training, the progression in the training sequence is a function of the trainee's performance based on simple linear relationships between the student's performance and the task difficulty. Adaptive training is also used as a means to increase student precision by reducing allowable error tolerance as student performance levels increase (Freedy, Lucaccini, & Lyman, 1967; Kelley & Wargo, 1967).

Overall, there have been few successful attempts to adapt training on the basis of entry characteristics. In addition, most successful adaptive CAI programs have been used in teaching lower order cognitive or perceptual-motor skills. An exception is the field of medicine, where promising CAI programs (mostly involving simulation) which teach skills involving higher order probabilistic inferences (e.g., diagnosis) have been written.

Problems of Adaptive CAI

Two problems are crucial to the development of adaptive CAI. The first is the problem of developing suitable models to describe student behavior. The second is the problem of optimizing instructional effectiveness on the basis of a description of the student in terms of model parameters. Solution of the former problem is clearly a prerequisite to the solution of the latter. Effective optimization techniques cannot compensate for inadequate or incorrect models of student behavior.

Sophisticated optimization techniques for maximizing instructional effectiveness have been used in several very elegant and highly adaptive CAI programs (e.g., Atkinson, 1972; Smallwood, 1971). However, these techniques have only been used for simple learning situations, which usually involve lower order cognitive skills such as memorizing lists of vocabulary words, because the optimization methods require a precisely stated learning model which predicts student response to alternate instructional options. As skills become more complex, it is less likely that simple mathematical learning models can be found.

AI and Adaptive Decision Models

A promising approach to adaptive CAI is the application of Artificial Intelligence (AI) techniques. Traditionally, AI techniques and theory have been concerned with problem-solving and decision-making tasks. These techniques are uniquely suitable for applications which involve unstructured environments (Nilsson, 1965; Slagle, 1971). Of particular interest are techniques which use trainable decision and classification networks. Such a technique is used as a basis for the CDDT system, since it provides unique capabilities to establish decision strategies through on-line observation of the decision process.

Work on adaptive decision-making is derived from the areas of behavioral decision research and AI experience with learning networks. The unique aspect of this approach is the capability to adjust model parameters on-line and change decision strategy accordingly. In essence, the learning system attempts to identify the decision process of the human operator in real time by: (a) successive observation of his actions; and (b) establishment of a model, i.e., an interim relationship between the input data and the output decision. Learning in this context refers to a training process for adjusting model parameters according to a criterion function. The object is to improve model performance as a function of experience, or to match the model characteristics to those of the operator.

Learning techniques have been used to model the decision strategy of the human operator and to identify the sources of cognitive constraints on the operator performing a dynamic prediction task (Rouse, 1972). Another example of an adaptive model of the human operator through real time parameter tracing has been reported by Gilstad and Fu (1970). Linear and piecewise-linear discriminant functions were used to classify system gains, errors and error rate. The decision boundaries for classification were determined through a process of on-line learning, observing operator performance and parameter adjustment. The specific model used was applicable only to very limited tasks, and merely illustrated the feasibility of the technique.

A unique advantage of using a learning system lies in its capability to act as a pattern classification mechanism. As such, it can be used to identify biases in operator decision policy as a response to classes or patterns in the input data (Tversky, Slovic, & Lichtenstein, 1972). In conventional Bayesian techniques, the pattern of events is decomposed into elementary data points. With the assumption of independence, the elementary data points are aggregated to revise the hypothesis. Effects of the data pattern do not influence the decision.

In dynamic decision making, however, the temporal and spatial nature of the data are highly significant. Since decision data appear as a pattern of individual events, it is reasonable to assume that the subject responds to the pattern as well as to the individual values. In fact, the pattern may contain the greater amount of information.

Classification of input patterns by the learning mechanism can be accomplished by programmed cognizance of such data features as: data with non-independent events; data with correlated events; data with events which continuously vary with time; the number of elements of decision data; and the rate of change in the data points.

Recently, new AI techniques have been introduced to take greater advantage of the unique adaptive capabilities of computers. Techniques such as natural language-understanding and pattern recognition have been used in CAI systems which are based on information structure representations of the subject matter (Carbonell, 1970; Hartley & Sleeman, 1973; Koffman & Blount, 1974; Brown, Burton, & Bell, 1974). These systems use network analysis of the structures to generate instructional sequences, thus the term "generative CAI."

The CDDT System

A promising approach for identifying the decision policy of the individual in dynamic processes is offered by the use of a trainable, multifunction decision mechanism which both models the Decision Maker (DM) and automates his or her decision functions (Freedy & Weltman, 1973). This approach offers both a framework for an operator model and an efficient means for handling the difficult problem of parameter identification in a stochastic environment.

Parameter adjustment is performed through "on-the-job decision tracking." That is, the decision network follows the decision policy of the DM and adjusts its parameters in order to make it behave like the operator. In essence, the decision network observes and acquires the decision policy of the DM. This includes identification of the process by which the DM maps classes of data or data patterns into diagnostic opinions about the environment, and the process by which component dimensions of utilities are aggregated into a single net utility.

The technique centers around the adjustment of an expected utility model. A maximum-likelihood model of real world model behavior is used to predict environment-state transitions and an expected utility model of decision maker behavior is employed to predict (and suggest) operator decisions. Both conditional probabilities of state transitions and the operator's utilities are estimated by the system. Currently the technique is being used in experimental investigation of the factors which influence optimal decision aiding in complex, realistic open intelligent gathering tasks (Freedy, Weisbrod, Davis, May, & Weltman, 1974; Freedy, May, Weisbrod, & Weltman, 1974). This approach forms the foundation for the CDDT system and is described in greater detail in Appendix A.

The CDDT system represents a departure from previous applications of AI techniques to CAI. Rather than using heuristic techniques of pattern recognition in the context of an information network representation

of the subject matter, the present system employs a learning system technique of pattern recognition within a framework of a decision model of the student and instructor. In this latter case, the subject matter is represented as a state space within which the student and instructor make decisions.

Decision Training in Maintenance

Selection

In the CDDT system adaptive sequential decision training is implemented within the context of electronic troubleshooting. The student's task is to find a circuit fault by making circuit measurements, replacing the malfunctioning part, and making final measurements to be able to declare that the device is repaired.

Electronic troubleshooting was chosen as the initial application of the adaptive decision methodology for several reasons. First, it is an ideal context in which to teach higher order skills involving judgment and probabilistic inference. Second, the troubleshooting task provides a firm basis for examining the applications of generalized concepts such as probability, values, and information gain to the training process. Third, electronic troubleshooting is an important skill in the military, to which CAI can be effectively applied. Finally, troubleshooting can be implemented on economical minicomputer systems because the problem is easily represented using either a state-space or a state-equation format. This means that the troubleshooting task can be formulated in a way which eliminates costly simulation of visual or tactile task elements. Accordingly, in the military electronic maintenance shop, as well as the classroom, a CAI terminal could provide the means for maintaining and sharpening skill levels during slack periods in the shop schedule. The CDDT system would also enable technicians to obtain practice in troubleshooting situations which are important but occur only seldomly or in emergencies.

Approach

The training given in the circuit fault diagnosis and repair task assumes that the student has a good basic background in electronics but that his experience with troubleshooting is limited. Such might be the case with a student who has recently completed advanced military electronics training but has not yet performed troubleshooting tasks in his first permanent duty assignment. This skill level can be assessed either in terms of previous training received or in terms of performance on an entering test of electronics and troubleshooting knowledge. It is assumed that the student has mastered the prerequisite laws of electricity, circuit component behavior, circuit subsytem, circuit diagrams, and the use of test equipment.

Training in the CDDT system occurs with certain restrictions on the extent of circuit simulation. The circuit has been restricted to include functional modules, rather than individual circuit components. The student interacts with a terminal display of the simulated circuit, thus he cannot make such troubleshooting observations as smelling faulty capacitors, looking for burned resistors, or touching overheated semiconductors. In addition, the measurement results are presented in a semi-interpreted form (high, normal, low, etc.), rather than as absolute readings (3.6 volts, 1.25 mA) so that the student need not refer to a table of normal circuit levels. These simplifications do not affect the inherent judgmental nature of the troubleshooting task.

2. THE CDDT SYSTEM

Expected Utility Decision Model

The CDDT system uses an Expected Utility (EU) model of the student and the instructor. In the student model, the EU principle is used as a basis for estimating the student's utilities for action outcomes. In the instructor model, the EU principle is used to generate suggested actions in response to student requests for assistance. The student's utilities for outcomes, in comparison with the expert's utilities for the outcomes, provide the framework for generating remedial instructions. The utilities provide a measure of the relative real world desirability or worth of an outcome.

Model Definition

The EU model is a prescriptive model which makes use of a criterion for determining the optimum choice among alternatives, assuming "rational behavior" (Edwards, 1962; Fishburn, 1964). The choice criterion employed is maximization of the individual expected relative utility as obtained by a weighted sum of individual utilities of consequences and their probability of occurrence. More specifically, the expected utility of an action is

$$EU_j = \sum_{i}^{n} P_{ij} U_{ij} \quad (1)$$

where

P_{ij} = probability that the i^{th} consequence in a set of n consequences will occur if action A_j is selected by the decision maker.

U_{ij} = relative utility of the i^{th} consequence of the j^{th} action.

Given a set of actions, utilities, and probabilities, the optimum choice can be determined according to the maximum EU principle by calculating the EU for each action and selecting the action with the highest EU.

The maximum EU principle has become a widely acceptable normative decision model for risky decision making (Luce & Raiffa, 1957; Krantz, Luce, Suppes, & Tversky, 1971). The work of Tversky (1967); Goodman, Saltzman, Edwards, and Krantz (1971); and others has indicated that the expert maximization principle provides a good first approximation for decision making under risk.

The EU model is used in the CDDT system as a basis for defining optimum strategies and as a structure for adaptive estimation of the student's utilities as inferred from his or her decision behavior. Using on-line adjustment of the student utility structure, an adaptive learning network estimates student utilities which can explain his or her actions by the criterion of maximization of EU. Thus the model continuously tracks the student's decision strategy as it changes during the course of training. In the instructor model, the EU model is used in real time as a criterion for recommending actions to the student in response to the student's requests for assistance. Off-line, the EU model is used to estimate the instructor's utilities.

EU Model Application

The EU model is used to define instructor and student choice behavior in selecting courses of action in a diagnosis and decision task. The probabilities of outcomes provide a measure of diagnostic progress, while the utilities provide a measure of the relative real world desirability or worth of an outcome. In the troubleshooting context, probabilities are associated with the likelihood of occurrence of measurement outcomes and circuit module faults as inferred from observed symptoms. The symptoms define the information available about the equipment at a given time. The utilities are associated with the worth of knowledge about certain action outcomes, and the contribution of this knowledge to determining technical circuit problems. In essence, the expected utility model defines the relative desirability of performing a certain measurement or part replacement under a given set of circuit symptoms. The model is expressed in terms of probabilities of obtaining an outcome--such as a certain measurement result--and the relative utility of the information about this measurement (discussed in following section). The instructor utilities are calculated using the adaptive utility estimation technique and then stored in a utility matrix. These instructor utilities are available during the training task to serve as a standard against which the estimated student utilities are compared. The aggregated probabilities of the instructor model are displayed to the student, thus justifying the use of the same set of probabilities in both the instructor model and the student model. Rather than using a static set of utilities as in the instructor model, the student utilities are dynamically adjusted throughout training using the adaptive, on-line utility adjustment subprogram.[1]

Information Gain

The expected utility model itself is insufficient to model dynamic decision behavior where the primary goal of the task is to gain information.

---

[1] Appendix A summarizes the techniques of dynamic utility estimation.

Such is the case with tasks of diagnosis, fault detection, intelligence gathering, and troubleshooting. Accordingly, it is necessary to introduce an information gain function to the EU equation. This has been approached as follows:

If we let a_{ij} equal the information gain resulting from obtaining outcome i of action j, the equation for EU is then written as follows:

$$EU_j = \sum_{i}^{n} a_{ij} \, P_{ij} \, U_{ij} \tag{2}$$

The most commonly used measure of information, I, is Shannon's formula (Shannon & Weaver, 1949):

$$I = -\sum_{k} P_k \, LOG_2 \, P_k \tag{3}$$

where P_k is the probability that the system is in state k. The question is how to use equation (3) to calculate a_{ij} in the context of the CDDT system.

It is possible to model a person's approach to troubleshooting in several ways. For example, initial experience showed that many people, when presented with the P_{ij}'s for a number of measurements j with outcomes i, tend to select the measurement with the greatest spread in the P_{ij}'s. Such a policy is modeled by the following information gain function.

$$a_j = -\sum_{i}^{n} P_{ij} \, LOG_2 \, P_{ij} \tag{4}$$

This represents a single information gain corresponding to all outcomes, i of action j, but since a_{ij} is inside the summation sign in Equation (2) we get:

$$a_{ij} = -P_{ij} \, LOG_2 \, P_{ij} \tag{5}$$

Another approach is taken by the troubleshooter who tries to eliminate the largest number of possible faults (regardless of probability or module in which the fault occurs) with each measurement. The following a models such a troubleshooter:

$$a_{ij} = f_i \, LOG_2 \, f_i \tag{6}$$

where, cconsidering previous measurements

$$f_i = \frac{\text{Faults associated with i}}{\text{All possible faults}} \tag{7}$$

Finally, someone who is more interested in verifying existing knowledge can be represented by the old EU model (Equation 1), in which:

$$a_{ij} = 1 \tag{8}$$

In the initial trials with the system, the information gain function represented by Equation (6) was used. The other information gain functions are also available at the experimenter's option.

System Organization

CDDT Structure

The structure of the CDDT system is illustrated in Figure 1. A training task is simulated for display to the student and for initial adjustment of the instructor model parameters. The instructor model includes probabilities of the occurrence of action outcomes which are displayed to the student. The instructor model also provides the standard of performance against which the student's behavior is evaluated and forms the basis of the training instructions. The student model provides the framework for dynamic on-line adjustment of the estimated student utilities. The set of instructional heuristics generate the instructions used to direct the student's troubleshooting behavior in the desired direction.

Individual components of the system are described in the following sections.

Task Simulator

The task simulator generates circuit faults and simulates the processes of checking symptoms, taking measurements, and replacing modules. Measurement results are simulated by reading them from a table of measurement results for each fault and measurement. The result of the requested measurement is displayed to the student in a semi-interpreted form (similar to the approach of Bond and Rigney, 1966) which eliminates the need for a manual of proper measurement values.

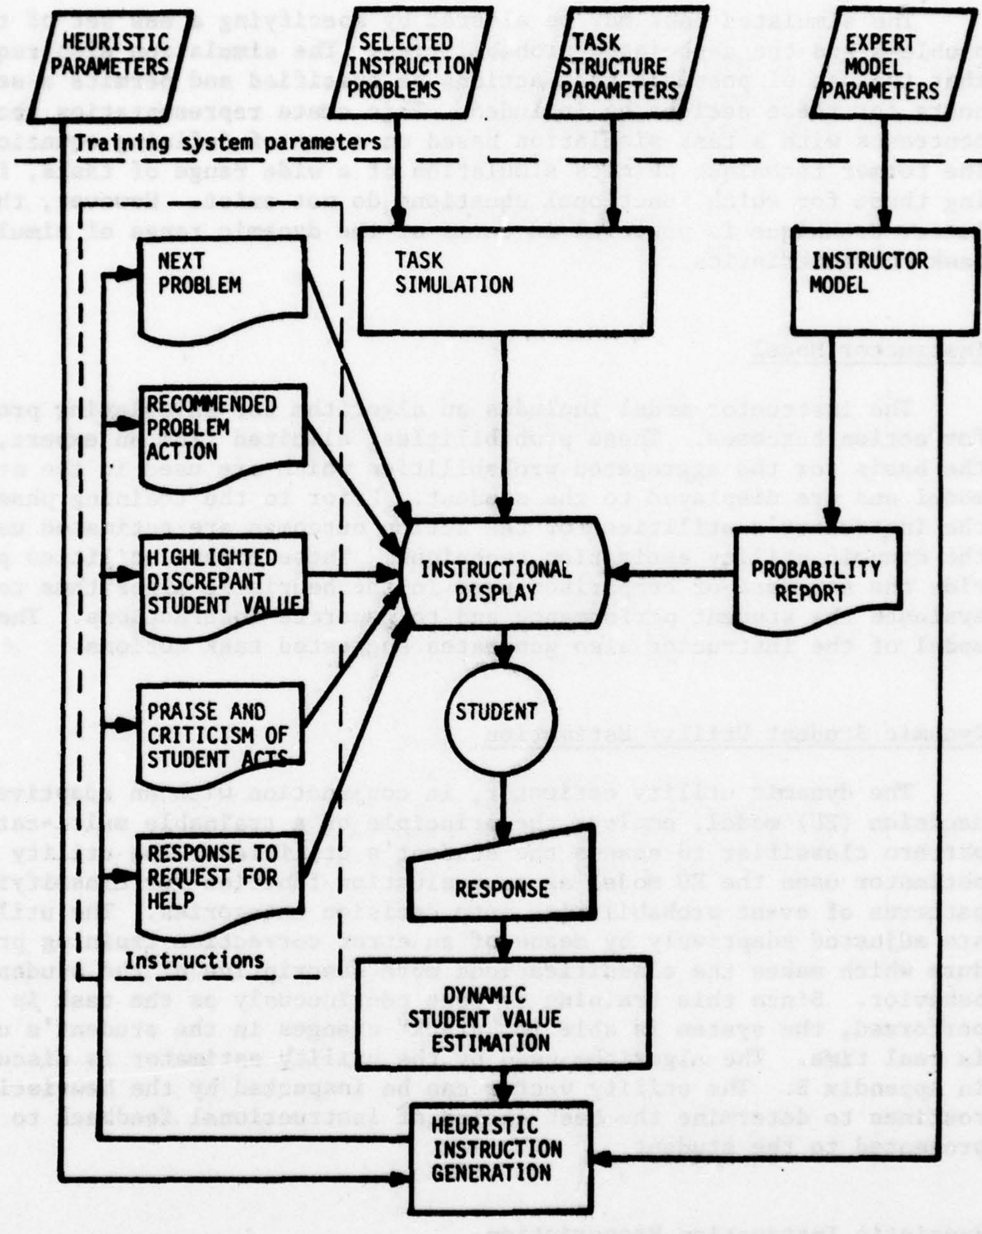

Figure 1. CDDT Functional Organization

The simulated task may be altered by specifying a new set of task problems and the associated probabilities. The simulation also requires that the set of possible task actions be specified and permits a set of costs for these actions be included. This state representation technique contrasts with a task simulation based on a set of defining equations. The former technique permits simulation of a wide range of tasks, including those for which functional equations do not exist. However, the latter technique is powerful in terms of the dynamic range of simulated task characteristics.

Instructor Model

The instructor model includes an algorithm for calculating probabilities for action outcomes. These probabilities, elicited from an expert, form the basis for the aggregated probabilities which are used in the student model and are displayed to the student. Prior to the training phase, the instructor's utilities for the action outcomes are estimated using the dynamic utility estimation technique. These expert utilities provide the standard of comparison used in the heuristic algorithms to evaluate the student performance and to generate instructions. The EU model of the instructor also generates suggested task actions.

Dynamic Student Utility Estimation

The dynamic utility estimator, in conjunction with an adaptive decision (EU) model, employs the principle of a trainable multi-category pattern classifier to assess the student's utilities. The utility estimator uses the EU model as an evaluation function for classifying patterns of event probabilities into decision categories. The utilities are adjusted adaptively by means of an error correction training procedure which makes the classifications more descriptive of the student's behavior. Since this training is done continuously as the task is being performed, the system is able to "track" changes in the student's utilities in real time. The algorithm used by the utility estimator is discussed in Appendix B. The utility vector can be inspected by the heuristics routines to determine the best choice of instructional feedback to be presented to the student.

Heuristic Instruction Prescription

Heuristic algorithms are used to find discrepancies between student and instructor utilities and to select instructional feedback aimed at reducing these discrepancies. The algorithms also evaluate the appropriateness of student actions. These algorithms are expandable, based on user experience.

The instructional heuristics consist of the criteria and the decision rules which select feedback, in the form of diagnostic messages and selected problems, on the basis of the student utilities. The criteria for instructional feedback can be classified into two types: (a) feedback which is based on a comparison of student utilities with those of the instructor, and (b) feedback which is based on analysis of the student's utilities alone. Both types of feedback are concerned with prescribing instruction for the student, based on the state of his estimated utilities as inferred from his decisions. Both the utilities and their temporal behavior in terms of convergence characteristics and rate of convergence are used in this form of feedback.

The extraction of prescriptive information from the student's utilities is the central focus in the CDDT system concept. In particular, the analysis of utilities provides a direct measure of decision-making consistency for the student, an indication of whether he or she approaches the correct utilities, and a measure of the rates at which he or she approaches the correct utilities. Discrepancies between the student utilities and those of the instructor model provide direct diagnostic information regarding student decision-strategies and areas where special instruction are required. The criterion for detecting discrepancies involves the computation of the difference between instructor and student utilities. A discrepancy is defined whenever the difference exceeds a certain threshold.

Actions for which discrepancies are detected form the basis for selecting the problem areas which are presented to the student. The problems are selected such that the relevant action which contributes to the solution of the problem involves the action space for which discrepancies have been identified. Thus, for a student utility that is too low relative to the instructor utility, the instructor would choose to obtain the measurement result in question, whereas the student would not choose to obtain that result. Conversely, if the student utility is too high, the instructor would not choose to obtain a result but the student would so choose.

This procedure of fault selection provides opportunities to reward appropriate measurement selections and to punish inappropriate selections. If the student responds correctly to the initial statements of utility discrepancies, the selected faults will assure opportunities to reinforce the appropriate measurement selections. If the student does not respond to the initial statements, the selected faults will cause him to continue to request measurements inappropriately and provide the opportunity to instruct the student within the context of the specific inappropriate behavior that has been exhibited. Several faults which are not directly related to the highlighted utility discrepancies may also be presented to avoid focusing attention solely on the specified measurement outcomes.

Requests for Help

In addition to the above, a "help" routine which uses the ideal response characteristics of the instructor model is available. This routine interprets a student's request for assistance and responds to student inquiries by: (a) listing potentially faulty modules; (b) suggesting an optimum action; and (c) suggesting an optimum action from the set of actions that the student is already considering. The first function involves checking the current outcome probabilities to determine which modules have a significant probability of being faulty at the current state of the diagnostic cycle. The second involves selecting the action, from among all possible actions, which has the highest expected utility according to the expert EU model. The third function involves selecting the action, from among the actions that the student is considering, which has the highest expected utility according to the expert EU model.

System Software

The CDDT software offers promise of providing an economical, efficient means of implementing CAI on a minicomputer. The system is designed for a time-sharing environment, thus permitting simultaneous training of several students. It is also context-free and can be readily modified since the subject matter is represented as a state space, rather than in terms of defining equations.

Functional Organization

The overall organization of the software system is shown in Figure 2. The major models of the system are described below. A number of other service routines, not discussed here, are shared by all modules.

Master System Scheduler. System control is centralized in the Master System Scheduler (MSS). MSS establishes and initializes all data areas, and schedules program flow. MSS has been designed to allow CDDT to be implemented in a time-sharing environment; i.e., to allow many independent students, each dealing with different circuit faults at various stages of instruction, to interact with the system simultaneously.

Graphics. The CAI display is generated and managed by a graphics programming package. Graphics software incorporated into the CDDT allows the MSS to modify the circuit display, for example, to highlight selected circuit modules, or nodes. Display elements can be modified off-line using a display compiler. Extensive modification of the display can be made with minimal programming effort.

Performance Monitor. These routines monitor the statistical behavior of the utilities for evaluation of system and student performance.

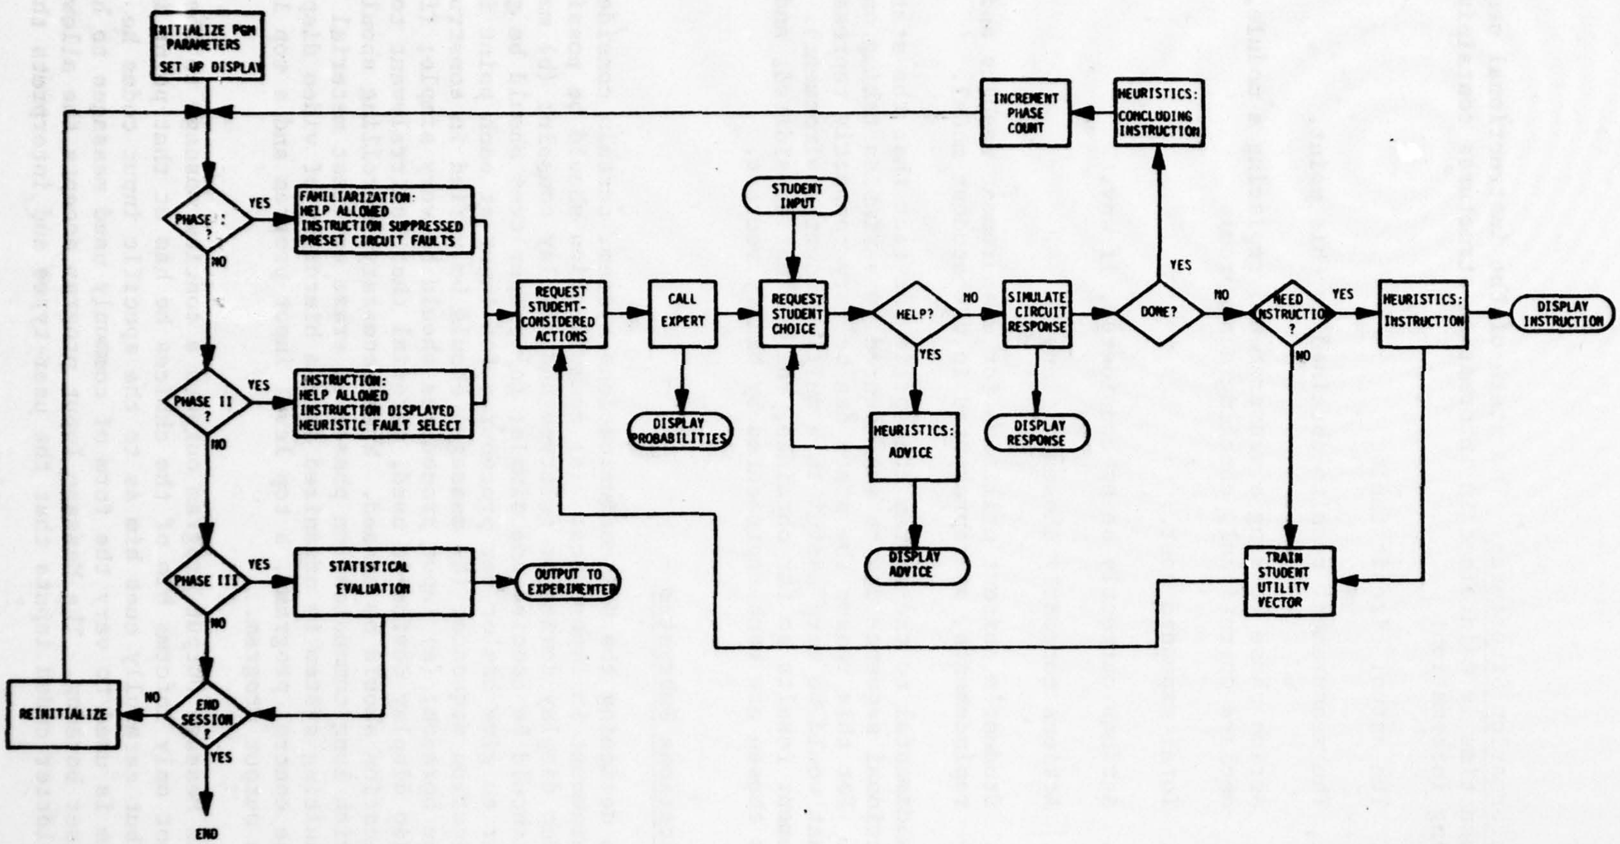

Figure 2. CDDT Software Functional Organization

Information Structures. The state of the instructional sequence at any given time is maintained in information structures containing the following information:

1. The actual circuit fault.

2. The measurement results obtained to this point.

3. Action phase (taking a measurement, replacing a module, help, declare operational, checking a symptom).

4. Total expended cost.

5. Actions currently being considered, if any.

6. Actions currently chosen, if any.

7. Student's current utilities for measurement results and module replacements, as represented in the student model.

Fundamental to the system concept is the fact that the state of the instructional sequence can be represented by variables taking on discrete values. For this reason the state can be very compactly represented (a fact that would be very useful in a multi-student environment). The measurement results so far obtained, the actions considered, and the actions chosen are each represented by binary vectors.

Communications Subsystem

In designing the CAI communications system, certain considerations were paramount in importance: (a) communication should be possible on the video display device or teletype and display console; (b) messages output should be concise and simple; (c) proper cues should be given to the user to give him or her procedures for input at each point in the communication sequence; (d) messages should be varied in construction to decrease boredom; (e) input procedures should be very simple; (f) when the video display device is used, material that is irrelevant to current communication should be erased. When necessary, scrolling should take place with long communication phases to erase earliest material first. The resulting system is organized into a hierarchy of video display and teletype control programs, a top level input program and a top level message output program.

The Message Output program outputs a concise message to the user which not only informs him of the choices he has at that point in the cycle, but carefully cues him as to the specific input codes he can use. A system is used to vary the form of commonly used messages to help avoid user boredom. The Message Input program accepts the allowed one- or two-letter coded inputs that the user types and interprets them. All

input as well as output is scrolled when necessary. When a new phase in the communication cycle begins, irrelevant previous material is erased from the communication area. The user is allowed to proceed at his or her own rate. New information is output only after the user pushes the transmit button indicating that he or she has read and understood the previous material and is ready to go on.

In the course of developing the CAI software, we found that it was easy to optimize the communication system by representing each phase and the associated messages graphically on a story board. The flow of communication could then be made simple and appropriate to the task and each message could be simplified and made consistent without communication conventions.

Training Environment

System Hardware

The CDDT system is implemented on an Interdata Model 70 minicomputer with 32K bytes of core memory.[1] Man/computer communication occurs on a teletype and on an Information Displays, Inc. IDIgraf graphic display terminal with 3K bytes of internal memory and direct memory access.

Instructor/Computer Interaction

The instructor (or experimenter) interacts with the CAI system primarily through the teletype. The system is designed to allow the instructor to modify, with a minimum of effort, the nature and complexity of the task environment, the decision model performance characteristics, and the structure of the student/computer interface.

The instructor controls the task environment by modifying the characteristics of the measurements and the circuit faults. He or she can modify the fault behavior of the circuit by changing the probability values, and can add new faults and measurements by making additional entries in the tables which define the information structures.

The performance characteristics of the decision model can be controlled by modifying (a) the initial utility values used by the adaptive EU model, (b) the learning rate of the utility estimator, and (c) the EU evaluation function used by both the model and the utility estimator. The easiest to modify are the initial values of the utility matrix, which are input by the instructor during program initialization. These initial values affect the behavior of the adaptive EU model and

---

[1]Commercial designations are used only for precision of description. Their use does not constitute endorsement by the Army or ARI.

the utility estimator, at least during the early stages of the training process. The learning rate of the utility estimator is controlled by a correction increment. This parameter affects the rate of convergence of the utility estimator and determines its sensitivity to changes in the operator's decision behavior. The size of the correction increment also affects the amount of variance which will result from inconsistent operator behavior. The most difficult method of controlling the decision model is modification of the expected utility function. This function, also used as a discriminant function by the utility estimator, is programmed into the system. Such modification of the EU function might be done, for example, if new types of task actions were required in the training system.

Student/Computer Interface

Figure 3 illustrates the student terminal display format. The display screen shows the problem circuit schematic, a list of possible measurements (keyed to points in the schematic), and the most recent instructional message. When a module is replaced, its outline is brightened as well. Interaction begins with a request for a student response (as exemplified in Figure 3). The student types the response and hits a "transmit" button; the response is entered into the system. A new message then appears in the message sector. Selected measurements are brightened on the display. Measurement results are displayed as they are obtained, and remain throughout the entire problem sequence.

Troubleshooting Task

Description

The student is introduced to the instructional system through a typewritten description of the troubleshooting task and the simulated problem circuit (see Appendix B). This same circuit diagram, as displayed on the student terminal, includes the modular blocks, the module interconnections, and the allowed measurements.

The present circuit is a power supply whose maintenance behavior has been thoroughly documented. The power supply has been divided into ten functional modules to resemble more closely a field repair task of an electronic system. Each module performs a well-defined role in the circuit, thus a malfunction in one module can be treated independently of other modules. This corresponds to more elaborate field-operational electronic systems which are designed with a modular approach, permitting rapid replacement of modules in the field, with detailed troubleshooting and component replacement occurring at a repair facility.

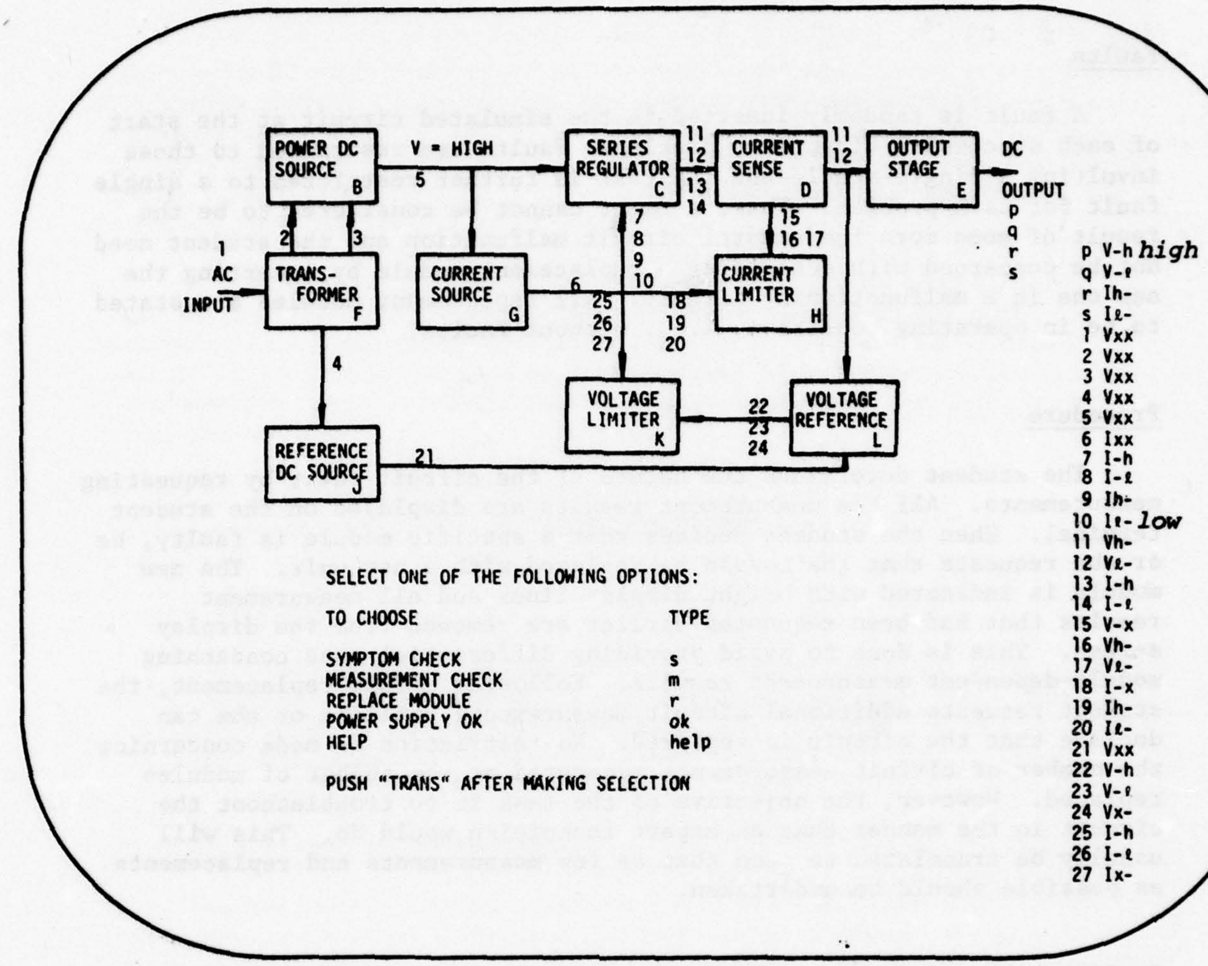

Figure 3. Student Terminal Display Format

Faults

A fault is randomly inserted in the simulated circuit at the start of each successive diagnostic problem. Faults are restricted to those involving a single module and the task is further restricted to a single fault for each problem. Thus, a fault cannot be considered to be the result of some more fundamental circuit malfunction and the student need not be concerned with destroying a replacement module by inserting the new one in a malfunctioning circuit. All replacement modules are stated to be in operating condition, i.e., without faults.

Procedure

The student determines the nature of the circuit fault by requesting measurements. All the measurement results are displayed on the student terminal. When the student decides that a specific module is faulty, he or she requests that the module be replaced with a new unit. The new module is indicated with bright display lines and all measurement results that had been requested earlier are removed from the display screen. This is done to avoid providing differential cues concerning module-dependent measurement results. Following module replacement, the student requests additional circuit measurements until he or she can declare that the circuit is repaired. No restriction is made concerning the number of circuit measurements requested or the number of modules replaced. However, the objective of the task is to troubleshoot the circuit in the manner that an expert technician would do. This will usually be translated to mean that as few measurements and replacements as possible should be undertaken.

Performance Motivation and Measurement

The troubleshooting task is presented within the context of a motivational structure that includes verbal commands but may also include a cost and payoff schedule. For many trainees, the introductory and instructional statements may well be sufficient to maintain cooperation and promote learning. However, long-term interest and performance improvement may prove to be difficult to maintain through instructions alone. In this regard, the cost/payoff structure can be imposed in the context of a game situation or the structure may be tied to a monetary or other incentive schedule. In the circuit troubleshooting task, the cost/payoff structure can be introduced within the context of a scenario that states the goal of the troubleshooting, presents the hypothetical background within which the task is being completed, the costs of taking the measurements and replacing modules, and the payoff for completing the task correctly.

Performance measures include the number of decisions and the time required to determine the fault, the expended cost of troubleshooting,

and the number of mistaken fault identifications. Further measures of system performance include number of requests for assistance and the number of perseverance errors.

Training Sequence

Program Phases

Troubleshooting problems are presented in two phases, familiarization and instruction.

Familiarization. A preselected or random set of problems is presented. No instructional (help) material is available. This phase provides a period for initial on-line adjustment of the estimated student utilities.

Instruction. Figure 4 illustrates the instructional flow diagram. Presentation of problems during the instruction phase is based primarily on the examination of the estimated student utilities. The instructional prescriptions that have been included to date provide a procedure for displaying to the student utilities for measurement outcomes that are most discrepant from the instructor utilities. In the present implementation, the two most discrepant utilities are selected and the student is instructed on the amount of discrepancy and what should be done to reduce the discrepancy. Subsequently, problems which emphasize faults related to these discrepancies are selected for presentation. Following completion of the selected fault problems, the two most discrepant utilities remaining are highlighted and the cycle is repeated. This instructional sequence is continued until the student and expert utility matrices correspond to the degree specified in the instructional objectives.

Additional machine responses not related to the examination of the utility matrices are also given during the instructional phase. These include such items as responses to specific requests for assistance and statements of allowable actions. For example, the student cannot declare the circuit to be repaired immediately following a module replacement, but must request further measurements to ascertain that the circuit is functioning normally. If such a declaration is made immediately following module replacement, an instruction will be given stating the appropriate action that is required.

Student Actions

Figure 5 illustrates the sequence of student actions and computer responses to these actions occurring within a single diagnosis problem.

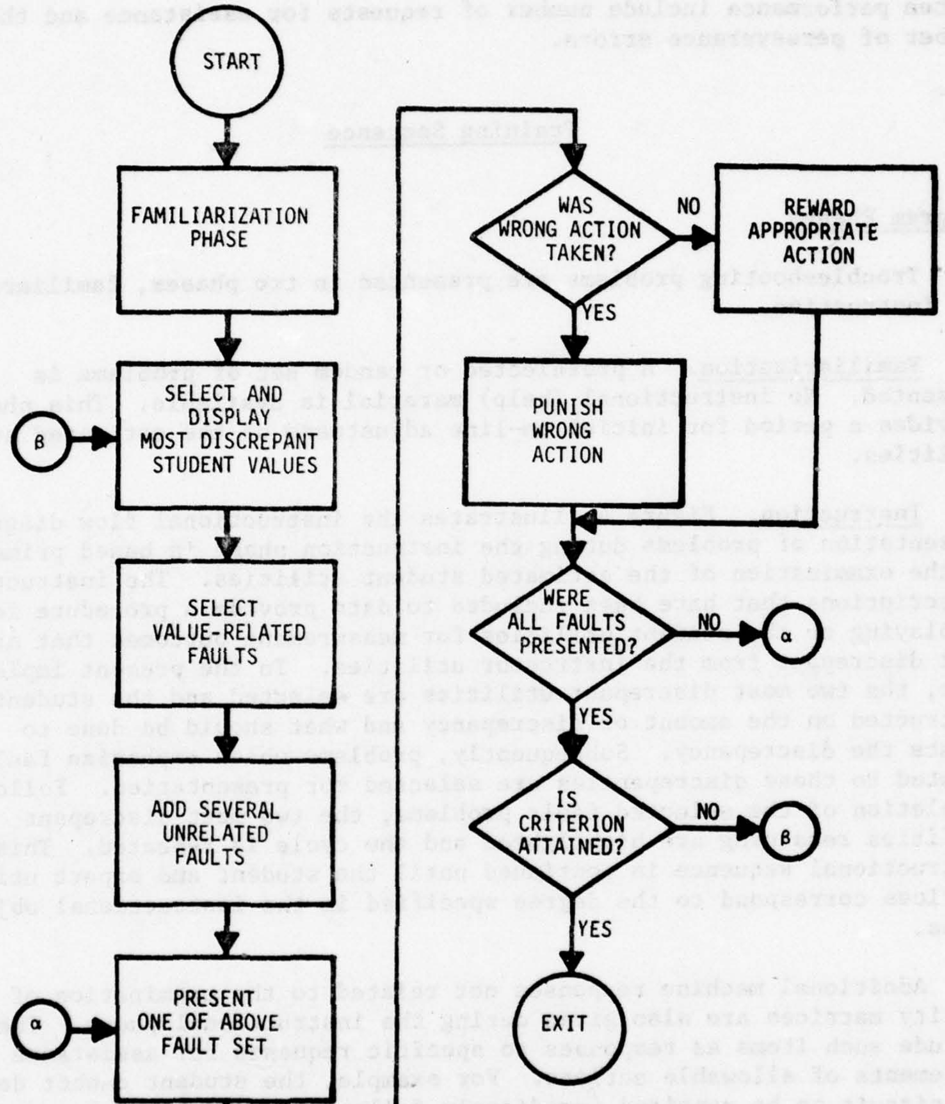

Figure 4. Fault Selection and Instruction Sequence

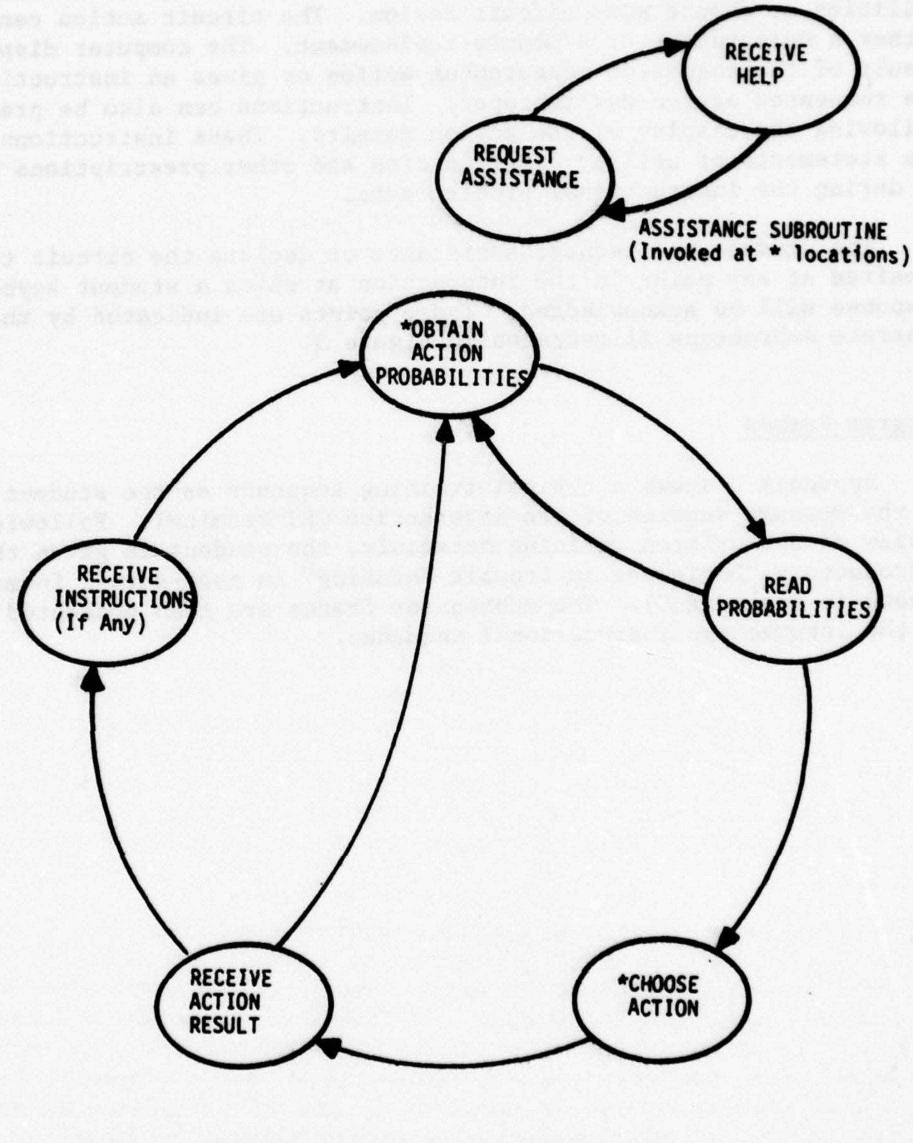

Figure 5. Student's Action Sequence

To begin, the student lists those circuit actions that he or she wishes to consider and the computer responds by listing the probability of obtaining each possible outcome for each of the actions listed. Following the probability presentation, the student can request further probabilities or choose some circuit action. The circuit action can be either a measurement or a module replacement. The computer displays the result of the requested measurement action or gives an instruction if the requested action was improper. Instructions can also be presented following the display of the action results. These instructions include the statements of utility discrepancies and other prescriptions generated as during the instructional problem sets.

The student may request assistance or declare the circuit to be repaired at any point in the interaction at which a student keyboard response will be acknowledged. These points are indicated by the separate subroutine illustrated in Figure 5.

Program Frames

Appendix D shows a typical training sequence as the student sees it in the message section of the interactive CRT terminal. Following review of the written training materials, the student is given the introductory "Decisions in Trouble Shooting" in page-scroll format (shown in Appendix C). The subsequent frames are then presented as part of the interactive instructional sequence.

3. RESULTS AND DISCUSSION

Accomplishments

The current year's efforts have fulfilled the initial program objectives, which were: (a) to apply an "intelligent" learning system to CAI and design a reliable adaptive training system; (b) to develop the instructional material and instructional sequence for system application in the context of an electronic troubleshooting task; and (c) to install the system on a minicomputer and demonstrate its operation.

Specifically, initial development of the student and instructor decision models in terms of structure and parameters has been finished. This involved conducting a decision analysis and mapping the troubleshooting task to an EU decision model. The result provided a structure for the implementation of the software modules and the adoption of the instructional criteria.

In addition, the first year of the research program involved the development and set-up of a basic operational CDDT system. Significant effort was directed toward the establishment of a system structure for the utilization of intelligent decision networks in CAI and the supporting subsystem to provide adaptive and individualized instructions. This included design and implementation of the equipment simulator, instructional environment, training sequence, graphic communication software, and the algorithms which provide instructional feedback on the basis of discrepancies between student and expert values. The effort culminated in demonstration of system operation and its functional integrity.

Contributions to CAI

The CDDT system represents a significant contribution to the general area of individualized, adaptive, computer-assisted instruction. Special features of the CDDT system include the following.

Adaptive Properties

CDDT supports an individualized approach to instruction within the context of mastery learning and specified levels of post-training competence. Specifically, CDDT:

(a) Provides an adaptive estimation of student (or expert) utilities within field-based decision making tasks.

(b) Permits individual student modes of learning and selects training problems based on the student's demonstrated capability.

(c) Considers individual strengths and weaknesses exhibited by the student in his decision utility structure; selection of future training problems is made in light of the student's particular needs and deficiencies.

(d) Presents training problems of increased complexity and generality as the student exhibits growing skills.

Diagnostic Capabilities

Generation of individual and group diagnostic information is facilitated by computerization. Tracking of progress and identification of problems becomes an integral part of the training sequence, including:

(a) Assessment of the adequacy of entry-level skills and knowledge of students initiating the CDDT program.

(b) Highlighting areas of individual student weakness during training for reference by student and instructor.

(c) Provision of statistical summaries that reflect the progress and current positions of individuals and the training group within the training program.

Flexibility

CDDT provides instructional flexibility in that adaptation to, or allowance for, the needs of individual students is possible. In the context of the total training situation, additional flexibility exists in that the CDDT system can be:

(a) Made available to students as class or work schedules permit.

(b) Easily modified through program replacement so that new training materials can be provided as desired.

(c) Useful under a number of training configurations.

Efficiency

Student instruction time is utilized efficiently because of the adaptive nature of CDDT and the active involvement required of the student. The system:

(a) Focuses on areas where the student makes the most errors; training is thus directed to those areas where the largest gains in skill are possible and needed.

(b) Provides a high density of responses and involvement during training problems. Selection of training problem difficulty levels on the basis of the student's demonstrated progress insures a work-load which is neither too light nor too heavy for the student's developing capability.

Preliminary Evaluation

Initial training system evaluation has been limited to date. It has focused primarily on the behavior of the utility estimates in the student model, and on the general "understandability" of the interaction. It was found that student utilities converged quickly under different student strategies, and that the model was consequently able to predict student behavior. This validated the estimation algorithms.

Early human factors evaluation of the interactive messages and instructional sequence led to a redesign into the format presented in this report. The current system has been judged easy to use by naive subjects. Both experienced and naive maintenance personnel have stated that the system captures the "feel" of an actual troubleshooting task. In particular, people without prior troubleshooting experience felt that the CDDT system is an ideal way to learn troubleshooting.

The CDDT system shows promise of being cost effective as well as educationally effective. The system design can be easily adapted for a multi-student environment. Since the instructional branching adapts to the individual student automatically, the cost of preparing lessons should be less than with conventional systems. The software is table driven. This means that most of the effort needed to change the instructional task consists of changing the tables, which is considerably simpler than reprogramming.

Recommended Future Work

Objectives

Recommended future work should include those tasks which are required to evaluate the system, define the optimum range of its major variables and expand the functional scope of adaptive instructions and student diagnosis. Specifically, the program should include the following objectives:

(a) To examine experimentally the effects of major instruction and feedback variables on system performance.

(b) To evaluate and optimize the human factors variables in the student interaction with the CDDT system.

(c) To expand the repertoire of heuristic instructional algorithms.

(d) To include the capability of diagnostic reports of student performance.

(e) To demonstrate and evaluate the instructional capabilities of the CDDT system.

Several major constituents are discussed further below.

Heuristics

The present CDDT system incorporates a limited set of algorithms which generate the instructional feedback. A major portion of the future effort should be directed toward the expansion of this set of heuristic algorithms. Such an expanded set should include methods of categorizing patterns of student utilities according to classes of requisite knowledge or hypotheses about troubleshooting. This classification procedure would facilitate the generation of instructions which are based on the clusters of utilities, rather than upon single utilities.

Utility Pattern Analysis

The adaptive prescription of instruction requires an analysis of utility patterns across a large set of action choices. Such patterns could then reflect weaknesses in specific technical areas rather than reflecting only an improper decision strategy or inability to select the proper action. Utility patterns are defined as a subset of the total utility space which contains utility discrepancies reflecing specific deficiencies in proficiency on the part of the student.

The utility pattern analysis extends the system from direct focus on student decision behavior to analysis and diagnosis of his actual state of knowledge. Furthermore, the use of utility patterns takes further advantage of AI techniques (in particular, the use pattern recognition techniques) to infer from student choices areas of lack of proficiency. The realization of this capability is most significant since it uses information that is available within the parameters of the student model and offers a significant increase in the capability of the CDDT system.

Diagnosis

A major element in the instructional effectiveness of a training system is its diagnostic capability. Accordingly, future work should include the further development of software for diagnosing student capabilities and weaknesses and for reporting these data in a form which is useful for improving the individual student's subsequent performance.

For those students who complete the familiarization and training phases, a diagnostic profile could be generated by the CDDT system. This profile would provide a statement of the student's performance in terms of the absolute levels of the dependent variables and in terms of the relation of the variable levels to the group performance previously established. The dependent variables in such a report would include the time required for the student's utilities to converge to stable levels, the specific utilities and patterns of utilities that have been identified, and the number of student requests for assistance.

References

Atkinson, R. C., Ingredients for a Theory of Instruction, American Psychologist, 27(10): 921-931, October 1972.

Atkinson, R. C. and Wilson, H. A. Computer-Associated Instruction. Science, 1968, 162, 73-77.

Baker, J. D., Computer Based Instruction System and Degree of Certainty, Newsletter, Greater Boston Chapter of the Society for Programmed Instruction, 1(4), 1965.

Bloom, B. S., Learning for Mastery, Evaluation Comment, (UCLA Center for the Study of Evaluation), 1968, 1(2).

Bond, N. A. and Rigney, J. W., Bayesian Aspects of Troubleshooting Behavior, Human Factors, 1966, 8(5), 377-383.

Brown, J. S., Burton, R. R., Bell, A. G. SOPHIE: A Sophisticated Instructional Environment for Teaching Electronic Troubleshooting. Cambridge, MA: Bolt Beranek and Newman, Inc., Technical Report BBN No. 2790, March, 1974.

Carbonell, J. R., AI in CAI: An Artificial-Intelligence Approach to Computer-Assisted Instruction, IEEE Transactions on Man-Machine Systems, 1970, MMS-11(4), 190-202.

Crawford, D. G. and Ragsdale, R. G., Individualized Quasi-Instructional Systems for the 70's, Working paper, Ontario Institute for Studies in Education, University of Toronto, July 1969.

Edwards, W., Dynamic Decision Theory and Probabilistic Information Processing, Human Factors, 1962, 4, 59-74.

Fishburn, P. C., Decision and Value Theory, New York: Wiley, 1964.

Freedy, A., Lucaccini, L. F., and Lyman, J., An Adaptive Approach to Optimum Switching Control of Artificial Limbs, Technical Report No. 67-47, Department of Engineering, UCLA, August 1967.

Freedy, A., May, D., Weisbrod, R., and Weltman, G., Adaptive Computer Aiding in Dynamic Decision Processes: Scenario Generation by Elicited Expert Probabilities. Woodland Hills, CA: Perceptronics, Inc., Technical Report PTR-1016-74-5(II), May 1974.

Freedy, A., Weisbrod, R., Davis, K., May, D., and Weltman, G., Adaptive Computer Aiding in Dynamic Decision Processes: Adaptive Decision Models and Dynamic Utility Estimation, Woodland Hills, CA: Perceptronics, Inc., Technical Report PTR-1016-74-5(I), May 1974.

Freedy, A. and Weltman, G., Adaptive Computer Aiding in Dynamic Decision Processes., Proceedings of the ARPA Conference on Decision Theory, RAND Corporation, Santa Monica, California, October 23-24, 1973.

Gilstad, D. W. and Fu, K. S., A Two-Dimensional Pattern Recognizing Adaptive Model of the Human Operator, Proceedings, Sixth Annual Conf. Manual Control, April 1970.

Glaser, R., Psychology and Instructional Technology, In R. Glaser (ed.) Training Research in Education, New York: Wiley, 1965.

Goodman, B., Saltzman, M., Edwards, W., and Krantz, D. M., When SEU is Irrelevant, Ann Arbor, MI: University of Michigan, Dept. of Psychology, Technical Report, 1971.

Groen, G. J. and Atkinson, R. C., Models for Optimizing the Learning Process, Psychological Bulletin, 1966, 66, 309-320.

Grubb, R. E., Learner-Controlled Statistics, In R. C. Atkinson and H. A. Wilson (eds.) Computer-Assisted Instruction, a Book of Readings. New York: Academic Press, 1969.

Grubb, R. E., CAI: Technical Aspects and Resources. In Computer Applications in Dental Education, San Francisco: U.S. Public Health Service, 1971.

Hartley, J. R. and Sleeman, D. H., Towards More Intelligent Teaching Systems, International Journal of Man-Machine Studies, 1973, 5, 215-236.

Kelley, C. R., What is Adaptive Training? Human Factors, 1969, 11, 547-556.

Kelley, C. R. and Wargo, M. J., Cross-Adaptive Operator Loading Tasks, Human Factors, 1967, 9, 395-404.

Koffman, E. B. and Blount, S. E., A Modular System for Generative CAI in Machine-Language Programming, IEEE Transactions on Systems, Man and Cybernetics, 1974, SMC-4 (4), 335-343.

Kopstein, F. F. and Seidel, R. J., The Computer as Adaptive Instructional Decision Maker, International Symposium on Man-Machine Systems, September 8-12, 1969.

Krantz, D. M., Luce, R. D., Suppes, P. and Tversky, A., Foundations of Measurement: Additive and Polynomial Representations, I. New York: Academic Press, 1971.

Lowes, A. L., Ellis, N. C., Norman, D. A., and Matheny, W. G., Improving Piloting Skills in Turbulent Air Using a Self-Adaptive Technique for a Digital Operational Flight Trainer, Orlando, FL: Naval Training Device Center, Technical Report NAVTRADEVCEN 67-C-0034-2, August 1968.

Luce, R. D. and Raiffa, H., Games and Decisions; Introduction and Critical Survey, New York: Wiley, 1957.

Lumsdaine, A. A. and Glaser, R. Teaching Machines and Programmed Learning. National Educ. Assoc., 1960.

Melaragno, R., A Study of Two Methods for Adapting Self-Instructional Materials to Individual Differences, System Development Corporation, Final Report, June 22, 1966.

Nilsson, N. J., Learning Machines, New York: McGraw Hill, 1965.

Rouse, W. B., Cognitive Sources of Suboptimal Human Prediction, Ph.D. Dissertation, MIT, Report DSR 70283-19, September, 1972.

Shannon, C. E. and Weaver, W., The Mathematical Theory of Communication, Urbana: University of Illinois Press, 1949.

Shuford, E. H., Cybernetic Testing and Degree of Certainty, paper presented at meeting of National Society for Programmed Instruction, Philadelphia, PA, May 5-8, 1965.

Shuford, E. H., Jr., and Massengill, H., The Relative Effectiveness of Five Instructional Strategies and Degree of Confidence, Evaluation Results, Interim Technical Report, the Shuford-Massengill Corporation, Lexington, MA, June 1967.

Slagle, J. R., Artificial Intelligence: The Heuristic Programming Approach, New York: McGraw-Hill, 1971.

Smallwood, R. D., A Decision Structure for Teaching Machines, Cambridge, MA: MIT Press, 1962.

Smallwood, R. D., The Analysis of Economic Teaching Strategies for a Simple Learning Model, J. of Math. Psychol., 8:285-301, 1971.

Stolurow, Lawrence M., Computer Assisted Instruction, in The Schools and the Challenge of Innovation, Committee for Economic Development, New York, 1969.

Suppes, P. and Morningstar, M., Computer-Assisted Instruction, Science, 1969, 166, 343-350.

Tversky, A., Additivity, Utility, and Subjective-Probability, Journal of Mathematical Psychology. 1967, 4, 175-202.

Tversky, A., Slovic, P., and Lichtenstein, S., Subjective Optimality: Conceptual Issues, Empirical Results and Decision Aids, in Application of Research on Cognition to Man-Machine System Design, ed. W. Edwards, Engineering Psychology Laboratory, University of Michigan, Report 010342-1-F, March 1972.

Zinn, Karl, An Evaluative Review of Uses of Computers in Instruction, Project CKUE, The University of Michigan Press, Ann Arbor, December 1970.

APPENDIX A

DYNAMIC UTILITY ADJUSTMENT

The dynamic utility estimation technique, developed by Perceptronics in the context of a decision aiding task (Freedy, Weisbrod, Davis, May, and Weltman, 1974), is based on the principle of a trainable multi-category pattern classifier. The utility estimator observes the operator's choices among R possible decision options available to him, viewing his decision making as a process of classifying patterns of event probabilities. The utility estimator then attempts to classify the event probability patterns by means of an expected utility evaluation, or discriminant, function. These classifications are compared with the operator's decisions and an adaptive error-correction training algorithm is used to adjust pattern weights, which correspond to utilities, whenever the classifications are incorrect. Thus, the value estimator "tracks" the operator's decision making and "learns" his utilities.

A multi-category pattern classifier (Nilsson, 1965) receives patterns of data and responds with a decision to classify each of the patterns in one of R categories. The classification is made on the basis of R linear discriminant (or evaluation) functions, each of which corresponds to one of the R categories. The discriminant functions are of the form

$$g_i(\overline{X}) = \overline{W}_i \cdot X \text{ for } i = 1, 2, \ldots, R \qquad \text{(A-1)}$$

where $\overline{X}$ is the pattern vector and $\overline{W}_i$ is a weight vector. The pattern classifier computes the value of each discriminant function and selects the category, i, such that

$$g_i(\overline{X}) > g_j(X) \qquad \text{(A-2)}$$

for all j = 1, 2, ..., R; $i \neq j$.

The adaptive error-correction training algorithm is very straightforward. Whenever the category selected by the pattern classifier, i, is different from the actual classification, k, the weights $\overline{W}_i$ are adjusted to reduce (punish) the value of $g_i(\overline{X})$ and the weights $\overline{W}_k$ are adjusted to increase (reward) the value of $g_k(\overline{X})$. Thus,

$$\overline{W}_i' = \overline{W}_i + d \cdot \overline{X} \qquad \text{(Reward)} \qquad \text{(A-3)}$$

$$\overline{W}'_i = \overline{W}_k - d \cdot \overline{X} \qquad \text{(Punish)} \qquad \text{(A-4)}$$

where d is the correction increment.

The dynamic utility estimator classifies pattern vectors

$$\overline{P} = [P_{1,1}, P_{1,2}, \ldots, P_{j,k}] \qquad \text{(A-5)}$$

whose components, P_{jk}, are the aggregated probabilities of the result, j, of action k. The discriminant functions are the expected utilities

$$EU_k = \sum_j P_{jk} \cdot U_{jk}$$

of the actions. The utility estimator computes the EU of each action and selects that action for which the EU is greatest. The estimator-selected action is compared with the action selected by the student and if they differ the appropriate utilities are rewarded (increased) or punished (decreased) by the training procedure. Thus the utilities are trained to characterize the operator's judgmental behavior, i.e., to make the utility estimator respond with the same decisions as the operator.

A fixed increment training rule is used to adjust the utilities. Whenever the action, d, selected by the utility estimator differs from the action, c, selected by the student, the utilities associated with the estimator action are punished and those associated with the student are rewarded:

$$U_j^{t+1}d = U_j^t d - (K \cdot P_j d) \qquad \text{(Punish)} \qquad \text{(A-6)}$$

$$U_j^{t+1}c = U_j^t c - (K \cdot P_j c) \qquad \text{(Reward)} \qquad \text{(A-7)}$$

The values at time t+1 are computed for all action results, j. The correction increment, k, is a constant which can be adjusted to give optimum convergence of the estimated utilities.

APPENDIX B

STUDENT HANDOUT

The Computerized Diagnosis and Decision Training (CDDT) system is a computer-based training system which is used to teach effective decision-making in electronic troubleshooting. You will be asked to locate circuit faults and to replace any malfunctioning circuit components. You locate the circuit malfunctions by making circuit measurements. The computer can help you locate the faults by giving you the probability of obtaining measurement outcomes, by suggesting what an expert troubleshooter would do at each point, and by showing you which measurements you overlook or which ones you use inappropriately. The object of the lesson is to learn to use the measurements so as to minimize the cost of troubleshooting. To do this you will learn to select measurements which will give you the most amount of information for their relative costs.

A regulated DC power supply is used as the sample electronic circuit. Rather than using the actual hardware device, a model of the circuit has been included in the computer system. You will be able to take circuit measurements by typing the measurement number on the computer keyboard. The simulated power supply will respond to the measurement and give a response which would be found in the actual power supply, given the malfunction which is to be located.

The simulated circuit is derived from the Heathkit model IP-28 regulated DC power supply with current limiting and voltage regulation features. The unit is designed to provide from 1 to 30 volts DC with a current limiting range from 10 milliamps to 1 amp. The IP-28 includes front-panel control of two voltage ranges and two current-limiting ranges with continuous potentiometer controls within the ranges. A meter also displays either current or voltage output and front-panel switches are provided for ON/OFF control of the AC input and the DC output. The circuit diagram of the IP-28 power supply is illustrated in Figure 6.

Because the objective of the lesson is for you to learn to make effective decisions concerning the choice of measurements, the circuit diagram has been simplified to emphasize the measurements, not the circuit components themselves. The circuit has been divided into a set of functional modules. Each module operates as a single unit and you need only replace a module rather than the individual circuit component. The subdivided circuit diagram is illustrated in Figure 7. This diagram also shows the circuit measurements which can be made. You may refer to this diagram during the time that you work with the computer system.

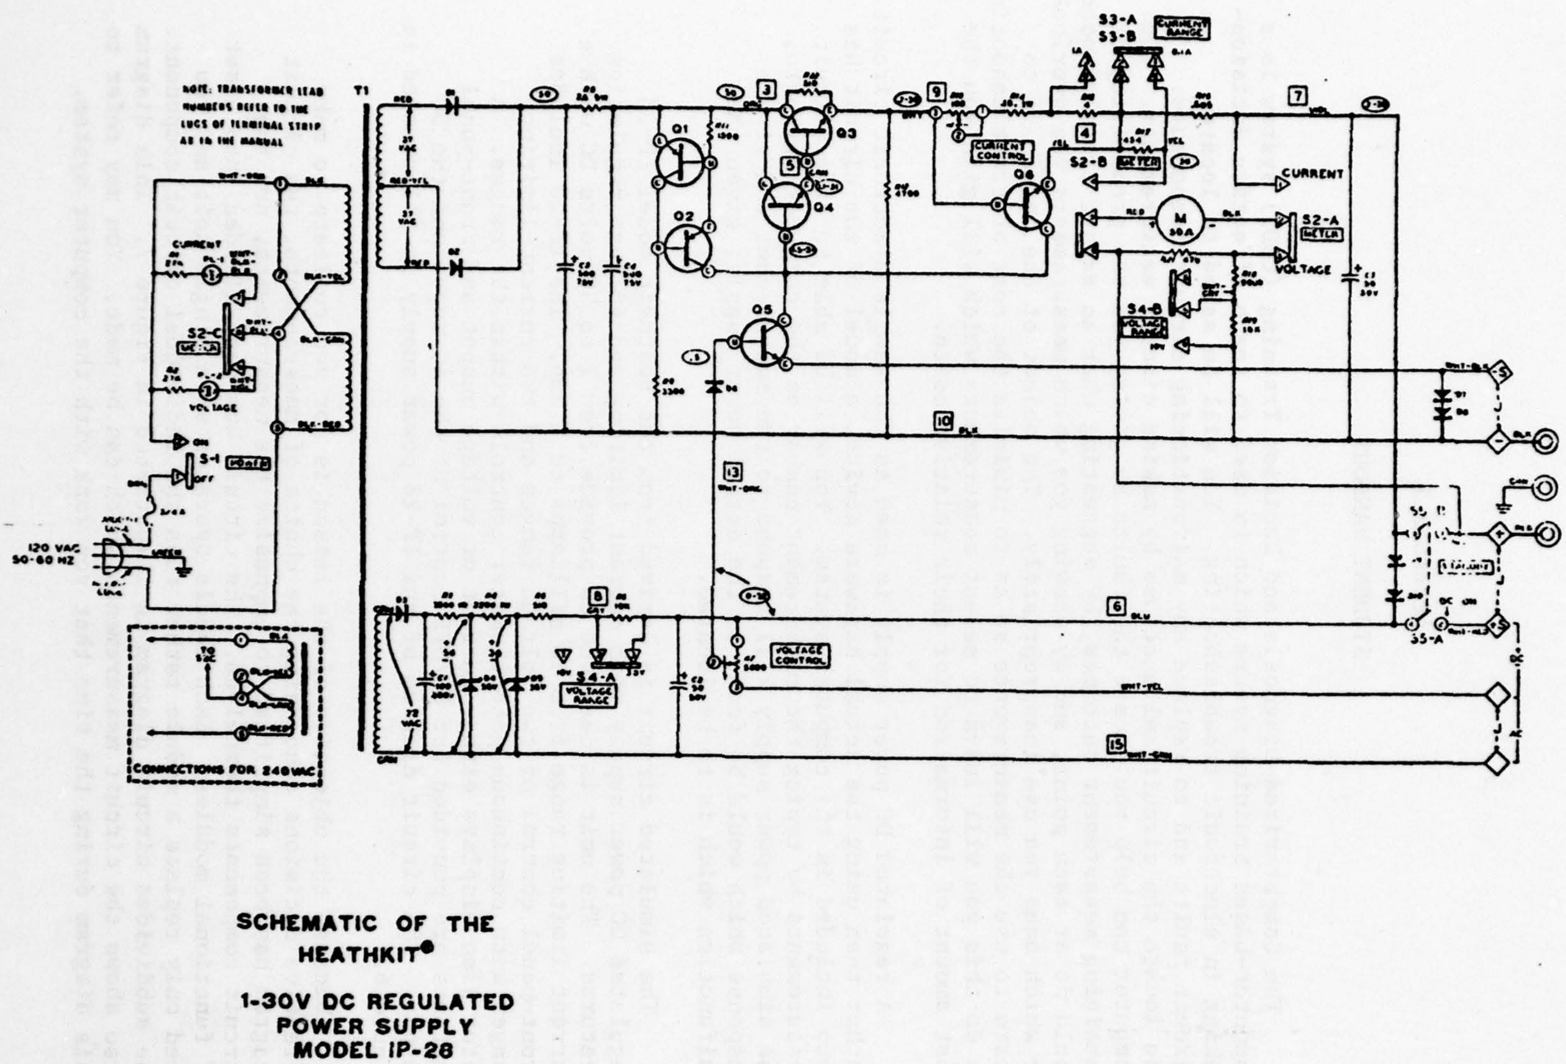

Figure 6. IP-28 Circuit Diagram

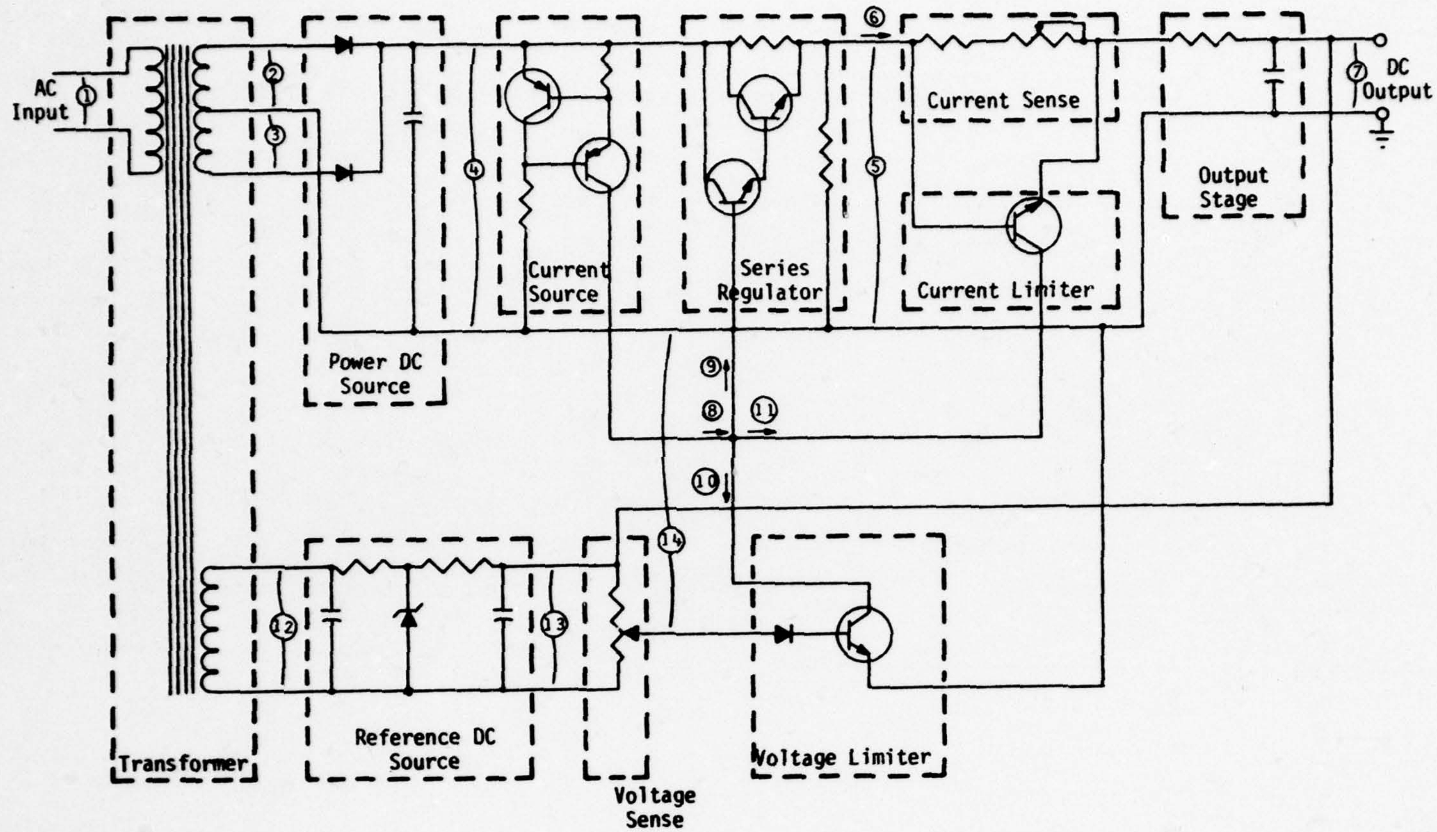

Figure 7. Simplified Circuit Diagram

APPENDIX C

DECISIONS IN TROUBLESHOOTING

THE FOLLOWING LESSON SEQUENCE IS DESIGNED TO HELP YOU MAKE EFFECTIVE DECISIONS WHEN YOU TROUBLESHOOT ELECTRONIC DEVICES. THE COMPUTER WILL PRESENT VARIOUS TROUBLESHOOTING PROBLEMS, YOUR JOB WILL BE TO LOCATE AND REPLACE ANY BAD MODULES, USING A SET OF AVAILABLE CHECK MEASUREMENTS. THE OBJECT IS TO TROUBLESHOOT THE CIRCUIT AT A MINIMUM COST. YOU CAN DO THIS BY SELECTING MEASUREMENTS THAT GIVE AS MUCH INFORMATION AS POSSIBLE FOR WHAT THEY COST. THE COMPUTER CAN HELP YOU BY SUGGESTING WHAT AN EXPERT TROUBLESHOOTER WOULD DO AT EACH POINT. IT CAN ALSO TELL YOU WHICH MEASUREMENTS YOU SEEM TO OVERLOOK OR WHICH YOU USE AT THE WRONG TIME.

THE FIRST FEW PROBLEMS WILL GIVE YOU PRACTICE ON THE COMPLETE TROUBLE-SHOOTING CYCLE. LATER IN THE LESSON, YOU WILL BE GIVEN CIRCUIT PROBLEMS WHICH HAVE BEEN PARTIALLY COMPLETED. THESE PROBLEMS GIVE YOU PRACTICE ON THOSE MEASUREMENTS THAT YOU HAVE NOT YET MASTERED. BY THE END OF THE LESSON SEQUENCE YOU WILL BE ABLE TO EXPERTLY DIAGNOSE THESE CIRCUIT PROBLEMS.

EACH PROBLEM WILL BE INTRODUCED AS A "DIAGNOSTIC PROBLEM." YOU CAN THINK OF IT AS SOMEONE BRINGING A DEVICE TO YOU WITH THE STATEMENT "IT DOESN'T WORK." YOU WILL NEED TO TAKE ENOUGH MEASUREMENTS TO DETERMINE HOW IT IS MALFUNCTIONING, REPLACE ANY BAD CIRCUIT MODULES, AND THEN STATE THAT THE CIRCUIT IS NOW OPERATIONAL. YOU MAY SOMETIMES FIND THAT THE CIRCUIT IS OK INITIALLY; BUT YOU MAY ASSUME THAT IF THE CIRCUIT DOESN'T WORK, ONLY ONE MODULE IS BAD.

WHEN YOU ARE READY TO PROCEED, PUSH THE "TRANS" BUTTON IN THE UPPER LEFT HAND CORNER OF THE KEYBOARD.

THIS DIAGRAM REPRESENTS THE FUNCTIONAL MODULES OF A REGULATED DC POWER SUPPLY WITH CURRENT LIMITING AND VOLTAGE REGULATION FEEDBACK LOOPS. THE NUMBERS IN THE DIAGRAM INDICATE THE CIRCUIT MEASUREMENT POINTS WHICH ARE SHOWN AT THE RIGHT. THE SYMPTOMS, P, G, R, AND S, ARE THE OUTPUT VOLTAGES AND CURRENTS.

THE FIRST STEP OF EACH PROBLEM IS TO LIST THE ACTIONS THAT YOU WOULD LIKE TO CONSIDER. THE COMPUTER WILL GIVE YOU THE PROBABILITY OF OBTAINING THE VARIOUS POSSIBLE RESULTS FOR THE ACTIONS YOU ARE CONSIDERING. YOU WILL THEN USE THESE PROBABILITIES, AND THE COST OF TAKING AN ACTION, TO DECIDE WHICH ACTION TO CHOOSE. YOU CAN CHECK THE OUTPUT SYMPTOMS, CHECK THE CURRENT AND VOLTAGE MEASUREMENTS AT THE NUMBERED CIRCUIT LOCATIONS, OR REPLACE ANY OF THE MODULES. WHEN YOU HAVE FIXED THE POWER SUPPLY YOU MUST SAY THAT THE CIRCUIT IS OK BEFORE THE NEXT PROBLEM WILL BE PRESENTED.

PUSH "TRANS" WHEN YOU ARE READY TO WORK ON THE FIRST PRACTICE PROBLEM.

APPENDIX D

STUDENT TRAINING SEQUENCE

	COMMENTS	MESSAGE FRAME
A.	The student is given a new problem. The message is acknowledged by pressing "TRANS."	DIAGNOSTIC PROBLEM THE POWER SUPPLY MAY HAVE A BAD MODULE. IF SO, YOUR JOB IS TO FIND THE BAD MODULE AND REPLACE IT. TO PROCEED, PUSH "TRANS."
B.	The student is asked to select an action category. Following the selection the student than sees frame C, D, F, or H depending on the action category chosen.	SELECT ONE OF THE FOLLOWING OPTIONS: TO CHOOSE — TYPE SYMPTOM CHECK — s MEASUREMENT CHECK — m REPLACE MODULE — r POWER SUPPLY OK — ok HELP — help PUSH "TRANS" AFTER MAKING SELECTION.
C.	Following the selection of action category "m," the student is asked to consider a number of plausible measurements. The computer responds with a table of probabilities similar to frame E.	WHAT MEASUREMENTS WOULD YOU LIKE TO CONSIDER? TYPE EACH NUMBER FOLLOWED BY "TRANS." PUSH "TRANS" AGAIN WHEN LIST IS COMPLETE.
D.	Following selection of action category "s" the student is asked to consider some symptom checks. The computer responds with frame E.	WHAT SYMPTOMS DO YOU WANT TO CONSIDER? TYPE EACH LETTER FOLLOWED BY "TRANS." PUSH "TRANS" AGAIN WHEN LIST IS COMPLETE.

E. The computer gives the probabilities of the outcomes for each considered measurement or symptom check. In this example, symptoms p, q, and s were chosen. After a symptom check or measurement is chosen, the result is displayed beside the column of measurements on the right of the screen. Then the student sees frame B again.

THIS TABLE GIVES THE PROBABILITY OF EACH RESULT AND THE COST OF THE CHECKS YOU ARE CONSIDERING

-----PROBABILITY OF READING-----

MEAS SYMP	NORML	NOT NORML	ZERO	VERY LOW	LOW	HIGH	FLOAT	COST
p	1	13	5	74	3	4	--	5
q	75		--	5	--	20	--	3
s	21	--	--	--	75	4	--	7

WHICH CHECK DO YOU CHOOSE?

CHOOSE A CHECK BY TYPING ONE NUMBER, ONE LETTER, OR "NONE" FOLLOWED BY "TRANS."

YOU CAN ALSO ASK FOR ASSISTANCE BY TYPING k FOLLOWED BY "TRANS."

F. Following selection of action category, "r," the student is asked to consider some possible bad modules.

WHAT MODULES ARE YOUR CANDIDATES FOR REPLACEMENT?

TYPE EACH MODULE LETTER, FOLLOWED BY "TRANS" WHEN YOUR LIST IS COMPLETE.

G. The failure probabilities are given for those modules being considered for replacement by the student. After a module replacement is selected, the appropriate module is brightened on the screen. The student then sees frame B again.

THIS TABLE GIVES THE PROBABILITIES THAT THE MODULES YOU ARE CONSIDERING ARE GOOD OR BAD:

MODULE	GOOD	BAD
A	80	20
C	60	40
F	60	40

WHICH MODULE DO YOU WANT TO REPLACE?

TYPE ONE LETTER OR "NONE" FOLLOWED BY "TRANS."

H.	The student has requested help. Following the student's selection the student sees the corresponding combination of frames I, J, or K.	HI, I AM YOUR INSTRUCTOR. HERE IS THE HELP I CAN GIVE YOU: 1. A LIST OF MODULES WHICH ARE MOST LIKELY TO BE BAD. 2. WHAT ACTION AN EXPERT WOULD TAKE AT THIS POINT. 3. THE ACTION THAT AN EXPERT WOULD SELECT FROM AMONG THE ACTIONS YOU ARE CONSIDERING. SELECT ONE OR MORE OF THE ABOVE OPTIONS BY TYPING ITS NUMBER FOLLOWED BY "TRANS."
I.	The student has selected option 1 of HELP.	FROM THE EVIDENCE AVAILABLE SO FAR, ONLY ONE OF THE FOLLOWING MODULES COULD BE FAILING: B G L PRESS "TRANS" TO CONTINUE.
J.	The student has selected option 2 of HELP.	THE EXPERT WOULD CHECK MEASUREMENT 15. PRESS "TRANS" TO CONTINUE.
K.	The student has selected option 3 of HELP.	FROM THE ACTIONS YOU ARE CONSIDERING THE EXPERT WOULD REPLACE MODULE D. PRESS "TRANS" TO CONTINUE.
L.	The student incorrectly declared the circuit operational. The student then sees frame B again.	TRY AGAIN, THE POWER SUPPLY DOES NOT WORK YET. PRESS "TRANS" TO PROCEED.
M.	The student correctly declared the circuit operational.	CONGRATULATIONS! YOU HAVE FIXED THE POWER SUPPLY. TO PROCEED, PUSH "TRANS."

N. The student's values are analyzed to provide a diagnosis of performance. This diagnosis is provided following each problem after the student repairs the circuit and declares the circuit operational.

HERE IS AN ANALYSIS OF YOUR PERFORMANCE UP TO NOW:

YOU ARE OVERVALUING THE FOLLOWING CHECKS. WHEN YOU CONSIDER USING THESE CHECKS, THINK CAREFULLY ABOUT HOW MUCH INFORMATION THEY REALLY PROVIDE IN THAT CASE:

P
Q
3
4
10

YOU ARE UNDERVALUING THE FOLLOWING CHECKS. YOU SHOULD CONSIDER THESE CHECKS WHEN THEY MAY PROVIDE USEFUL INFORMATION:

R
S
8
26

PUSH "TRANS" TO CONTINUE.

O. A partially completed problem is generated to give the student practice where it is most needed. The student next sees frame B.

DIAGNOSTIC PROBLEM:

THIS PROBLEM HAS BEEN PARTIALLY COMPLETED. YOUR JOB IS TO FINISH THE DIAGNOSIS AND TO REPLACE ANY BAD MODULE.

THE LIST AT THE RIGHT GIVES THE RESULTS OF THE CHECKS THAT HAVE ALREADY BEEN TAKEN. YOU SHOULD CONSIDER THESE RESULTS BEFORE CHOOSING ACTIONS.

PUSH "TRANS" TO CONTINUE.